Juniper® Validated Design

JVDE Test Report Brief: 3-Stage Data Center Design with Juniper Apstra and VMware NSX-T

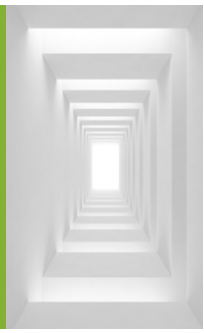

Introduction

This test report brief contains qualification test report data for the 3-Stage Data Center Design with Juniper Apstra and VMware NSX-T Juniper Validated Design Extension (JVDE). This qualification includes validation of the blueprint deployment, incremental configuration push, checking telemetry and analytics, validating data, and traffic flow. If you have questions about this Juniper Validated Design, contact your Juniper Networks representative.

The design is based on an ERB (Type 2 and Type 5) EVPN/VXLAN fabric with spine, server leaf, border leaf, and NSX-T Edge Gateway devices.

In the same way that server virtualization programmatically creates and manages virtual machines, NSX-T Data Center network virtualization programmatically creates and manages virtual networks. NSX-T Data Center works by implementing three separate but integrated planes: management, control, and data. These planes are implemented as a set of processes, modules, and agents residing on two types of nodes: NSX Manager and transport nodes.

The Juniper Data Center Fabric is viewed as an external network, which is defined as a physical network or VLAN not managed by NSX-T Data Center. You can link the NSX logical network or overlay network to an external network through a Tier-0 router or NSX Edge.

The goal is to ensure the design is well-documented and will produce a reliable, predictable deployment for your requirement.

The qualification objectives include validation of blueprint deployment, device upgrade, incremental config pushes/provisioning, Telemetry/Analytics checking, failure mode analysis, and verification of host traffic. Hardware models and software release for each device role will be explicitly listed.

Test Topology

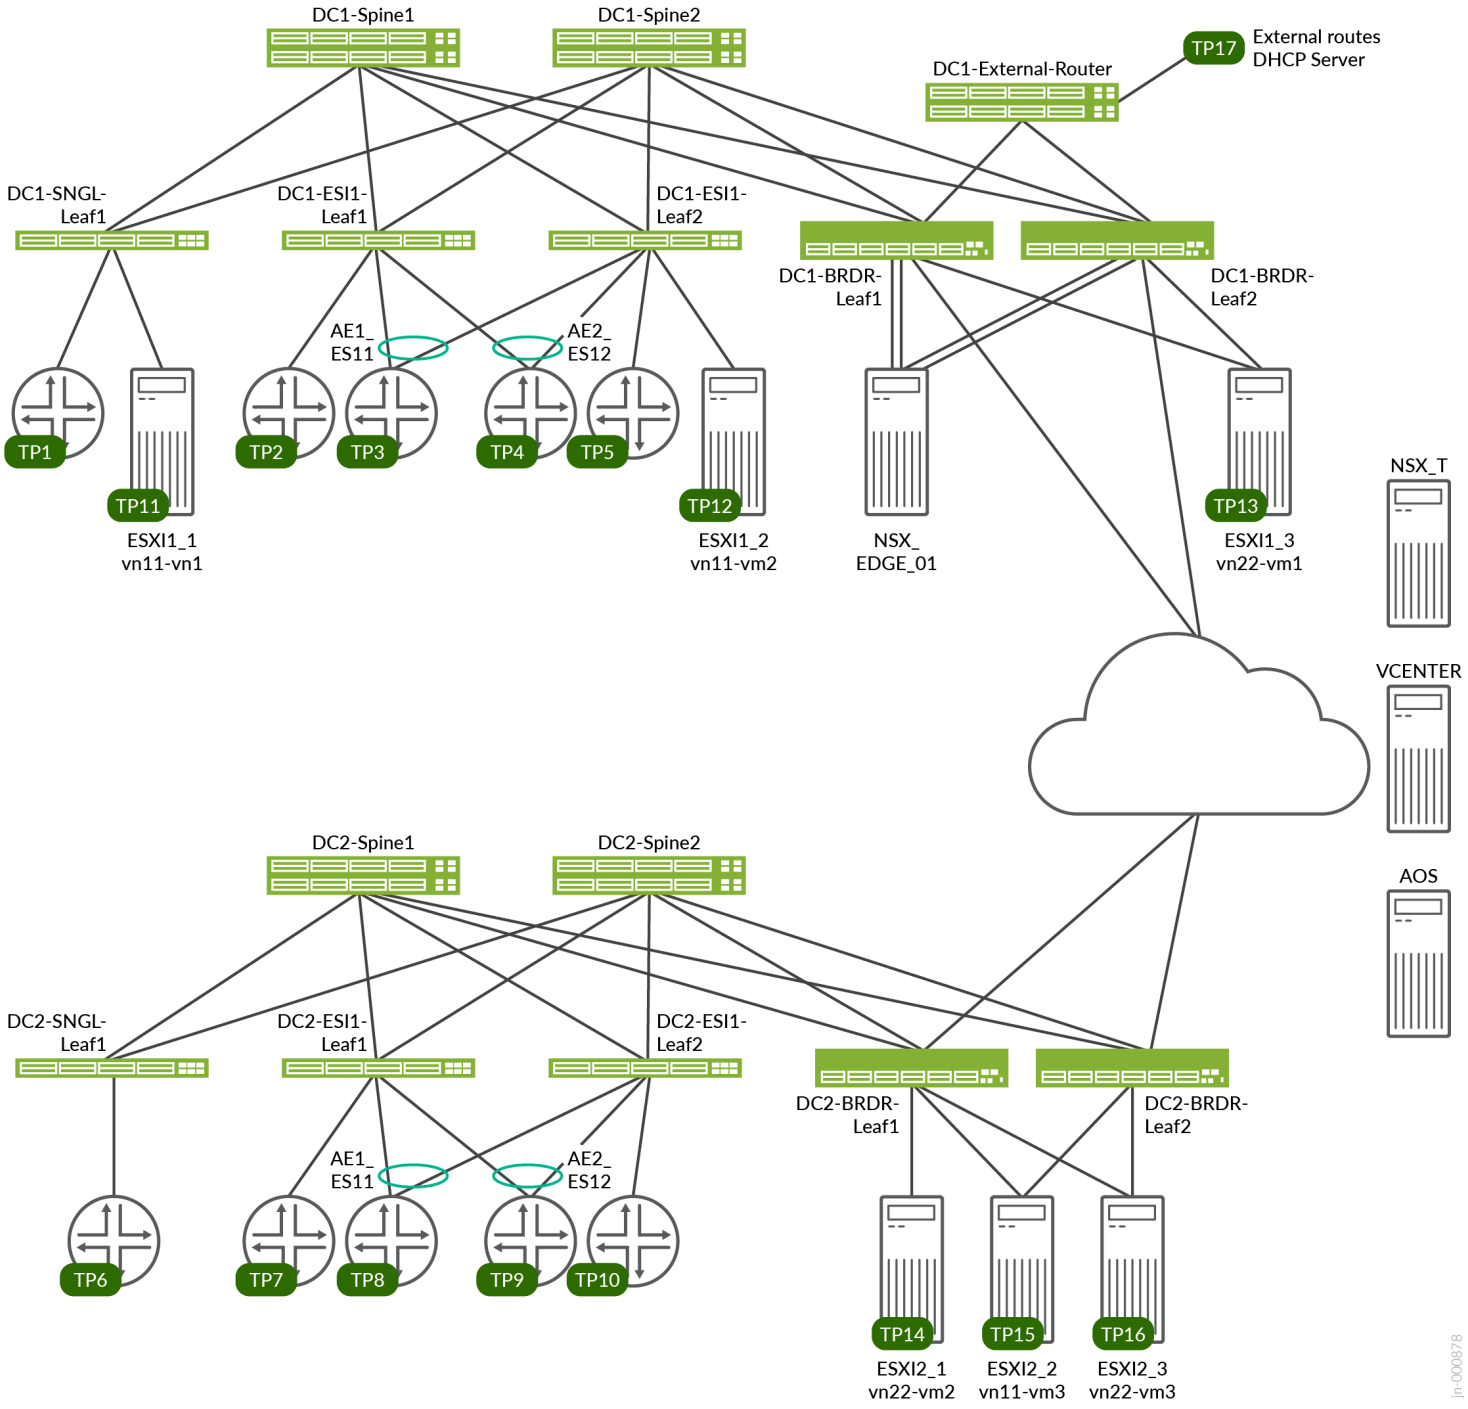

Figure 1: Reference Topology

Platforms Tested

The following table shows the platforms that were tested with this JVD.

Table 1: Tested Platforms

Role	Platform	OS
DC1-SNGL-LEAF1	QFX5120-48Y-8C	Junos OS Release 22.2R3-S3
DC1-ESI1-LEAF1	QFX5120-48Y-8C | QFX5110-48S-4C | EX4400-MP	Junos OS Release 22.2R3-S3
DC1-ESI1-LEAF2	QFX5120-48Y-8C | QFX5110-48S-4C | EX4400-MP	Junos OS Release 22.2R3-S3
DC1-BRDR-LEAF1	QFX10002-36Q | QFX5130-32cd | QFX5700 | ACX7100-48L | ACX7100-32c | PTX10001-36MR	Junos OS Release 22.2R3-S3
DC1-BRDR-LEAF2	QFX10002-36Q | QFX5130-32cd | QFX5700 | ACX7100-48L | ACX7100-32c | PTX10001-36MR	Junos OS Release 22.2R3-S3
DC1-SPINE1	QFX5220-32cd | QFX5120-32cd	Junos OS Release 22.2R3-S3
DC1-SPINE2	QFX5220-32cd | QFX5120-32cd	Junos OS Release 22.2R3-S3
TOR_SWITCH	QFX10002-36Q	Junos OS Release 22.2R3-S3
DC1-EXTERNAL-Router	MX304	Junos OS Release 22.2R3-S3
AOS	AOS	4.2.1

Version Qualification History

This JVD Extension has been qualified in Junos OS Release 22.2R3-S3 and Apstra AOS 4.2.1.

Scale and Performance Data

This document may contain key performance indexes (KPIs) used in solution validation. Validated KPIs are multi-dimensional and reflect our observations in customer networks or reasonably represent solution capabilities. These numbers do not indicate the maximum scale and performance of individual tested devices. For uni-dimensional data on individual SKUs, contact your Juniper Networks representatives.

The Juniper JVD team continuously strives to enhance solution capabilities. Consequently, solution KPIs may change without prior notice. Always refer to the latest JVDE test report for up-to-date solution KPIs. For the latest comprehensive test report, contact your Juniper Networks representative.

Table 2: Scaling Numbers

Feature	Tested Scale
DC1SLEAF1_VN_Count	500
DC1SLEAF1_VLAN_Count	500
DC1SLEAF1_IRB Count	500
DC1SLEAF1_local_mac_ip_host_entries	2500

Feature	Tested Scale
DC1SLEAF1_local_ndp_host_entries	200
DC1ESILEAF1_VN Count	500
DC1ESILEAF1_VLAN Count	1500
DC1ESILEAF1_IRB Count	500
DC1ESI1LEAF1_local_mac_ip_host_entries	7500
DC1ESI1LEAF1_local_ndp_host_entries	600
DC1ESILEAF2_VN Count	500
DC1ESILEAF2_VLAN Count	1500
DC1ESILEAF2_IRB Count	500
DC1ESI1LEAF2_local_mac_ip_host_entries	7500
DC1ESI1LEAF2_local_ndp_host_entries	600
DC1 total MAC IP count	17500
DC1 total NDP entry count	1400
VNI per leaf node	500
VTEP per leaf node	6
ESI per leaf node	4
BGP Routing Table per leaf node	363000
EVPN Table size per leaf node	44000

NOTE: These scale numbers are NOT device maximums and only reference the scale when these multidimensional test cases are performed.

Table 3: Performance Numbers

Event	Performance
Multihomed access link failure	Traffic recovery time < 750msec
Leaf-to-spine link failure	Traffic recovery time < 500msec
Dual-homed node reboot	Traffic loss impact < 500msec
Single node BGP protocol flap	Traffic loss impact < 150msec
Global Mac Learning initialization time for 20k entries	< 10 seconds

High Level Features Tested

The following table shows the features that were tested with this JVDE.

Table 4: Features

Feature	Node	Description
Single-homed Access Links	DC1-SNGL-LEAF1 DC1-ESI1-LEAF1 DC1-ESI1-LEAF2	Up to 2,000 VLANs per access interface, distributed between Blue and Red VRF. Access hosts use up to 10 MAC/IP, and 5 NDP/IPv6 entries per VLAN. DHCP client and server emulation on selected VLANs.
Multi-homed Access Links	DC1-ESI1-LEAF1 DC1-ESI1-LEAF2	Two AE bundles with ESI and LACP shared between LEAF1 and LEAF2. Up to 2,000 VLANs per access AE bundle, distributed between Blue and Red VRF. Access hosts use up to 10 MAC/IP, and 5 NDP/IPv6 entries per VLAN. DHCP client and server emulation on selected VLANs.
Border Leaf Pair	DC1-BRD-LEAF1 DC1-BRD-LEAF2	Apstra Border Leaf Rack Type with links to Generic System. BGP peering from each VRF to the Generic System to provide connectivity between VRFs, external DHCP server, and external hosts.
EBGP underlay and overlay	All Fabric Devices	Default routing profile for the Apstra Blueprint.
IP ECMP with fast-reroute	All Fabric Devices	Equal Traffic Distribution at all multipath points (i.e., Leaf-to-Spine1/Spine2), tested in steady state and with link/protocol/device/process flapping.
BFD	All Fabric Devices	Underlay and Overlay BFD with BGP at timers of both 1 sec and 3 sec respectively (default AOS setting) validated in steady state and with link/protocol/device/process flapping.
MAC-VRF	All Leaf Devices	Apstra default routing-instance type for Layer2. Single instance per leaf, with 1 VXLAN VNI per VLAN, and VLAN-aware service-type.
Layer 3 IRB interface	All Leaf Devices	One L3 IRB address per VLAN/VNI. Same address/MAC used on all leaves per VLAN/VNI. 9000 byte MTU.
EP-style Access interface	All Leaf Devices	Apstra default VLAN configuration style.
Type 2 routing	All Fabric Devices	Default Apstra L2 route announcement type.
Type 5 routing	All Fabric Devices	Default Apstra ip-prefix-routes announcement type.
DHCP Relay	All Leaf Devices	Per-VRF DHCP forwarding and relay-source override profile.
LLDP	All Devices	Apstra default link discovery mechanism.
Duplicate MAC detection	All Leaf Devices	Apstra default MAC-VRF setting.
BGP Graceful Restart	All Devices	All BGP peers.
NSX-T Uplink Profiles for edge and host nodes	Border Leaves	Border leaf peering with NSX Edge node to provide connectivity between NSX transport zones and to Apstra fabric hosts.

Event Testing

The following table shows the events that were tested with this JVDE.

Table 5: Tested Events

Test	Details
Create Agent profiles for managed devices	Apstra Deployment Step as documented in JVD

Test	Details
Create Pristine Configurations and acknowledge managed devices	Apstra Deployment Step as documented in JVD
Create a Logical Device, interface Map, and device profile for leaf devices	Apstra Deployment Step as documented in JVD
Create a Logical Device, interface Map, device profile for spine devices	Apstra Deployment Step as documented in JVD
Create a Logical Device, Interface Map, Device profile for external router and emulated servers	Apstra Deployment Step as documented in JVD
Create a rack type for ESI leaf pair + 4 servers	Apstra Deployment Step as documented in JVD
Create a rack type for single leaf	Apstra Deployment Step as documented in JVD
Create a rack type for border leaf	Apstra Deployment Step as documented in JVD
Create rack templates for single and ESI leaves	Apstra Deployment Step as documented in JVD
Create a rack template for border leaves	Apstra Deployment Step as documented in JVD
Create a blueprint with the template	Apstra Deployment Step as documented in JVD
Assign all resource pools for the blueprint	Apstra Deployment Step as documented in JVD
Assign Interface Maps to managed devices	Apstra Deployment Step as documented in JVD
Check Cabling Map	Apstra Deployment Step as documented in JVD
Commit the blueprint	Apstra Deployment Step as documented in JVD
Configure Overlay Network routing zones	Apstra Deployment Step as documented in JVD
Assign EVPN loopbacks	Apstra Deployment Step as documented in JVD
Create Virtual Networks	Apstra Deployment Step as documented in JVD
Commit Overlay Updates	Apstra Deployment Step as documented in JVD. Verify successful establishment of the control plane and verify no anomalies reported.
Prepare NSX-T elements	NSX Deployment
Build NSX and Junos OS communication via Apstra	NSX Deployment
Verify NSX-T Geneve tunnels are Up	NSX Deployment and control plane validation through Apstra DC Fabric.
Verify Overlay Connectivity	Validate forwarding plane establishment with emulated host test traffic.
Apstra Device State Changes	Undeploy/drain/set pristine configuration of leaf and spine nodes. Validate redeployment success and re-establishment of control and forwarding planes.
Remove and Add Tenants	Modify virtual networks and connectivity templates for selected VLANs and check for removal/addition with test traffic. Expected traffic loss when tenant is removed, and restoration of traffic once tenants are re-added.
Reboot Devices	Ensure minimal traffic loss when a redundant node is rebooted, and that control and forwarding plane are restored with ECMP when the node reboot is completed. Reboot NSX Edge Node and validate NSX Transport Zone connectivity is restored.
Server Link Failure	Interface down/up, laser on/off at access layer for single-homed link. Only traffic to/from the connected hosts will be impacted.
Multihomed link failure	Interface down/up, laser on/off at access layer with minimal traffic loss for links which are redundant. AE member links and complete AE bundles are flapped. Only traffic to/from the connected hosts will be impacted.

Test	Details
Leaf to Spine Link failure	Interface down/up, laser on/off at leaf-spine layer. Minimal traffic loss as all leaves are dual-homed.
Process Restart	Ensure minimal traffic loss and full recovery when various Junos OS processes are killed/restarted.
MAC Move	MAC/IP host moves within VLAN to adjacent leaf devices. vMotion move of vCenter VM to adjacent leaf devices.
Deactivate BGP on leaf	Ensure minimal traffic loss when various leaf and spine BGP sessions are deactivated. Full traffic restoration when sessions are restored.
Reset DHCP Bindings	DHCP release and renew.
Longevity Tests	Extended traffic run for 8 hours.
Extended Negative Testing	Continual loop of process restarts, protocol, and interface flaps to ensure stability and resilience of control and data plane.
Fabric Device Upgrade from Apstra	Junos OS Image changes performed through Apstra on all managed devices. Verify control and forward planes are functional after system upgrade.

Traffic Profiles

The following table shows the traffic profiles that were tested with this JVD.

Table 6: Tested Traffic Profiles

Traffic Path	Type	Load	Packet Size
dc1sleaf1_red_to_allred	Intra_VRF	1000pps	random 256-1024 bytes
dc1sleaf1_blue_to_allblue	Intra_VRF	1000pps	random 256-1024 bytes
dc1sleaf1_red_to_allblue	Inter_VRF	1000pps	random 256-1024 bytes
dc1sleaf1_red_v6_to_allred_v6	Intra_VRF	1000pps	random 256-1024 bytes
dc1esi1leaf1_red_v6_to_allred_v6	Intra_VRF	1000pps	random 256-1024 bytes
dc1esi1leaf1_red_to_allred	Intra_VRF	1000pps	random 256-1024 bytes
dc1esi1leaf1_blue_to_allblue	Intra_VRF	1000pps	random 256-1024 bytes
dc1esi1leaf1_red_to_allblue	Inter_VRF	1000pps	random 256-1024 bytes
dc1esi1leaf2_red_to_allred	Intra_VRF	1000pps	random 256-1024 bytes
dc1esi1leaf2_blue_to_allblue	Intra_VRF	1000pps	random 256-1024 bytes
dc1esi1leaf2_blue_to_allred	Inter_VRF	1000pps	random 256-1024 bytes
dc1esi1leaf2_blue_v6_to_allblue_v6	Intra_VRF	1000pps	random 256-1024 bytes
dc1esi1leaf2_sngl_red_to_allred	Intra_VRF	1000pps	random 256-1024 bytes
dc1esi1leaf2_sngl_blue_to_allblue	Intra_VRF	1000pps	random 256-1024 bytes
dc1esi1leaf2_sngl_blue_to_allred	Inter_VRF	1000pps	random 256-1024 bytes
dc1esi1leaf2_sngl_blue_v6-to_allblue_v6	Intra_VRF	1000pps	random 256-1024 bytes
External_routes_to_all_blue	External	1000pps	random 256-1024 bytes
External_routes_to_all_red	External	1000pps	random 256-1024 bytes

Known Limitations

When setting MAC-VRF instance on the QFX5120 and QFX5110 leaf devices, the following CLI command is required: **set forwarding-options evpn-vxlan shared-tunnels**. Once Apstra applies this command, the device will require a reboot. Apstra reports this as an anomaly until the device is manually rebooted.

When using QFX10002 as a border leaf and DHCP relay is required, you must set the no-snooping parameter through an Apstra configlet. For example: **set routing-instances blue forwarding-options dhcp-relay no-snoop**.

Corporate and Sales Headquarters
Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, CA 94089 USA
Phone: 888.JUNIPER (888.586.4737)
or +1.408.745.2000
Fax: +1.408.745.2100
www.juniper.net

APAC and EMEA Headquarters
Juniper Networks International B.V.
Boeing Avenue 240
1119 PZ Schiphol-Rijk
Amsterdam, The Netherlands
Phone: +31.207.125.700
Fax: +31.207.125.701